



[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: ☒ The ACM Digital Library ☐ The Guide
+estimate +cardinality +"data skew" +query +optimization for
SEARCH



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used estimate cardinality data
skew query optimization formula equation

Found 33 of 186,958

Sort results by
Display results

[Save results to a Binder](#)
 [Search Tips](#)
☐ Open results in a new window

Try an [Advanced Search](#)
Try this search in [The ACM Guide](#)

Results 1 - 20 of 33

Result page: [1](#) [2](#) [next](#)

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Paper session DB-8 \(databases\): query optimisation: Towards estimating the number of distinct value combinations for a set of attributes](#)



Xiaohui Yu, Calisto Zuzarte, Kenneth C. Sevcik
October 2005 **Proceedings of the 14th ACM international conference on Information and knowledge management CIKM '05**

Publisher: ACM Press

Full text available: [pdf\(269.74 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Accurately and efficiently estimating the number of distinct values for some attribute(s) or sets of attributes in a data set is of critical importance to many database operations, such as query optimization and approximation query answering. Previous work has focused on the estimation of the number of distinct values for a single attribute and most existing work adopts a data sampling approach. This paper addresses the equally important issue of estimating the number of distinct value combinati ...

Keywords: cardinality estimation, relational database

2 [Integration of spatial join algorithms for processing multiple inputs](#)



Nikos Mamoulis, Dimitris Papadias
June 1999 **ACM SIGMOD Record , Proceedings of the 1999 ACM SIGMOD international conference on Management of data SIGMOD '99**, Volume 28 Issue 2

Publisher: ACM Press

Full text available: [pdf\(1.66 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Several techniques that compute the join between two spatial datasets have been proposed during the last decade. Among these methods, some consider existing indices for the joined inputs, while others treat datasets with no index, providing solutions for the case where at least one input comes as an intermediate result of another database operator. In this paper we analyze previous work on spatial joins and propose a novel algorithm, called slot index spatial join (SISJ), t ...

Keywords: query optimization, spatial joins, spatial query processing

3 [Research sessions: statistics: A bi-level Bernoulli scheme for database sampling](#)



Peter J. Haas, Christian König

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available:  [pdf\(344.80 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#)

Current database sampling methods give the user insufficient control when processing ISO-style sampling queries. To address this problem, we provide a bi-level Bernoulli sampling scheme that combines the row-level and page-level sampling methods currently used in most commercial systems. By adjusting the parameters of the method, the user can systematically trade off processing speed and statistical precision---the appropriate choice of parameter settings becomes a query optimization problem. We ...

4 Fast algorithms for universal quantification in large databases



Goetz Graefe, Richard L. Cole

June 1995 **ACM Transactions on Database Systems (TODS)**, Volume 20 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(3.51 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Universal quantification is not supported directly in most database systems despite the fact that it adds significant power to a system's query processing and inference capabilities, in particular for the analysis of many-to-many relationships and of set-valued attributes. One of the main reasons for this omission has been that universal quantification algorithms and their performance have not been explored for large databases. In this article, we describe and compare three known algorithms ...


5 Space efficient bitmap indexing



Nick Koudas

November 2000 **Proceedings of the ninth international conference on Information and knowledge management**

Publisher: ACM Press

Full text available:  [pdf\(268.42 KB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

6 Advanced SQL modeling in RDBMS



Andrew Witkowski, Srikanth Bellamkonda, Tolga Bozkaya, Nathan Folkert, Abhinav Gupta, John Haydu, Lei Sheng, Sankar Subramanian

March 2005 **ACM Transactions on Database Systems (TODS)**, Volume 30 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(279.06 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Commercial relational database systems lack support for complex business modeling. ANSI SQL cannot treat relations as multidimensional arrays and define multiple, interrelated formulas over them, operations which are needed for business modeling. Relational OLAP (ROLAP) applications have to perform such tasks using joins, SQL Window Functions, complex CASE expressions, and the GROUP BY operator simulating the pivot operation. The designated place in SQL for calculations is the SELECT clause, whi ...

Keywords: Excel, OLAP, analytic computations, spreadsheet


7 Searching in metric spaces with user-defined and approximate distances



Paolo Ciaccia, Marco Patella

December 2002 **ACM Transactions on Database Systems (TODS)**, Volume 27 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(555.89 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Novel database applications, such as multimedia, data mining, e-commerce, and many others, make intensive use of similarity queries in order to retrieve the objects that better fit a user request. Since the effectiveness of such queries improves when the user is allowed to personalize the similarity criterion according to which database objects are evaluated and ranked, the development of access methods able to efficiently support user-defined similarity queries becomes a basic requirement. In t ...

Keywords: Distance metrics, user-defined queries

8 XML indexing and compression: Efficient processing of joins on set-valued attributes



Nikos Mamoulis

June 2003

Proceedings of the 2003 ACM SIGMOD international conference on Management of data

Publisher: ACM Press

Full text available: pdf(678.13 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citings](#), [index terms](#)

Object-oriented and object-relational DBMS support set valued attributes, which are a natural and concise way to model complex information. However, there has been limited research to-date on the evaluation of query operators that apply on sets. In this paper we study the join of two relations on their set-valued attributes. Various join types are considered, namely the set containment, set equality, and set overlap joins. We show that the inverted file, a powerful index for selection queries, c ...

9 Fast joins using join indices

Zhe Li, Kenneth A. Ross

April 1999

The VLDB Journal — The International Journal on Very Large Data Bases, Volume 8 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(263.06 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Two new algorithms, "Jive join" and "Slam join," are proposed for computing the join of two relations using a join index. The algorithms are duals: Jive join range-partitions input relation tuple ids and then processes each partition, while Slam join forms ordered runs of input relation tuple ids and then merges the results. Both algorithms make a single sequential pass through each input relation, in addition to one pass through the join index and two passes through a te ...

Keywords: Decision support systems, Query processing

10 Generalized multidimensional data mapping and query processing



Rui Zhang, Panos Kalnis, Beng Chin Ooi, Kian-Lee Tan

September 2005

ACM Transactions on Database Systems (TODS), Volume 30 Issue 3

Publisher: ACM Press

Full text available: pdf(689.08 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Multidimensional data points can be mapped to one-dimensional space to exploit single dimensional indexing structures such as the B^{plus}_t -tree. In this article we present a Generalized structure for data Mapping and query Processing (GiMP), which supports extensible mapping methods and query processing. GiMP can be easily customized to behave like many competent indexing mechanisms for multi-dimensional indexing, such as the UB-Tree, the Pyramid technique, the iMinMax, and the iDistan ...

Keywords: Indexing, data mapping, efficiency

11 Query evaluation techniques for large databases

Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Publisher: ACM Press

Full text available: pdf(9.37 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#),
[index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

12 Research sessions: query progress: Estimating progress of execution for SQL queries

Surajit Chaudhuri, Vivek Narasayya, Ravishankar Ramamurthy

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available: pdf(201.38 KB)

Additional Information: [full citation](#), [abstract](#), [references](#)

Today's database systems provide little feedback to the user/DBA on how much of a SQL query's execution has been completed. For long running queries, such feedback can be very useful, for example, to help decide whether the query should be terminated or allowed to run to completion. Although the above requirement is easy to express, developing a robust indicator of progress for query execution is challenging. In this paper, we study the above problem and present techniques that can form the basi ...

13 Research papers: OLAP: Efficient computation of multiple group by queries

Zhimin Chen, Vivek Narasayya

June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available: pdf(371.92 KB)

Additional Information: [full citation](#), [abstract](#), [references](#)

Data analysts need to understand the quality of data in the warehouse. This is often done by issuing many Group By queries on the sets of columns of interest. Since the volume of data in these warehouses can be large, and tables in a data warehouse often contain many columns, this analysis typically requires executing a large number of Group By queries, which can be expensive. We show that the performance of today's database systems for such data analysis is inadequate. We also show that the pro ...

14 Research sessions: compression: Wavelet synopses with error guarantees

Minos Garofalakis, Phillip B. Gibbons

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02**



Publisher: ACM Press

Full text available: pdf(1.44 MB)



Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#),
[index terms](#)

Recent work has demonstrated the effectiveness of the wavelet decomposition in reducing large amounts of data to compact sets of wavelet coefficients (termed "wavelet synopses") that can be used to provide fast and reasonably accurate approximate answers to queries. A major criticism of such techniques is that unlike, for example, random sampling, conventional wavelet synopses do not provide informative error guarantees on the accuracy of individual approximate answers. In fact, as this paper de ...

15 Processing queries for first-few answers

 Roberto J. Bayardo, Daniel P. Miranker
November 1996 **Proceedings of the fifth international conference on Information and knowledge management**
Publisher: ACM Press
Full text available:  [pdf\(894.67 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)


16 Probabilistic wavelet synopses

 Minos Garofalakis, Phillip B. Gibbons
March 2004 **ACM Transactions on Database Systems (TODS)**, Volume 29 Issue 1
Publisher: ACM Press
Full text available:  [pdf\(668.62 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Recent work has demonstrated the effectiveness of the wavelet decomposition in reducing large amounts of data to compact sets of wavelet coefficients (termed "wavelet synopses") that can be used to provide fast and reasonably accurate approximate query answers. A major shortcoming of these existing wavelet techniques is that the quality of the approximate answers they provide varies widely, even for identical queries on nearly identical values in distinct parts of the data. As a result, users ha ...

Keywords: Wavelets, approximate query processing, data synopses, randomized rounding



17 On applying hash filters to improving the execution of multi-join queries

Ming-Syan Chen, Hui-I Hsiao, Philip S. Yu
May 1997 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 6 Issue 2
Publisher: Springer-Verlag New York, Inc.
Full text available:  [pdf\(318.66 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

In this paper, we explore an approach of interleaving a bushy execution tree with hash filters to improve the execution of multi-join queries. Similar to semi-joins in distributed query processing, hash filters can be applied to eliminate non-matching tuples from joining relations before the execution of a join, thus reducing the join cost. Note that hash filters built in different execution stages of a bushy tree can have different costs and effects. The effect of hash filters is evaluat ed fir ...



Keywords: Bushy trees, Hash filters, Parallel query processing, Sort-merge joins

18 Industrial sessions: big data: Automating physical database design in a parallel database

 Jun Rao, Chun Zhang, Nimrod Megiddo, Guy Lohman
June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02**
Publisher: ACM Press
Full text available:  [pdf\(1.38 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



Physical database design is important for query performance in a shared-nothing parallel database system, in which data is horizontally partitioned among multiple independent nodes. We seek to automate the process of data partitioning. Given a workload of SQL statements, we seek to determine automatically how to partition the base data across multiple nodes to achieve overall optimal (or close to optimal) performance for that workload. Previous attempts use heuristic rules to make those decision ...

19 On parallel execution of multiple pipelined hash joins

 Hui-I Hsiao, Ming-Syan Chen, Philip S. Yu
May 1994 **ACM SIGMOD Record , Proceedings of the 1994 ACM SIGMOD international conference on Management of data SIGMOD '94**, Volume 23 Issue 2
Publisher: ACM Press
Full text available:  [pdf\(1.24 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this paper we study parallel execution of multiple pipelined hash joins. Specifically, we deal with two issues, processor allocation and the use of hash filters, to improve parallel execution of hash joins. We first present a scheme to transform a bushy execution tree to an allocation tree, where each node denotes a pipeline. Then, processors are allocated to the nodes in the allocation tree based on the concept of synchronous execution time such that inner relations (i.e., hash tables) ...

20 Implementing sorting in database systems

 Goetz Graefe
September 2006 **ACM Computing Surveys (CSUR)**, Volume 38 Issue 3
Publisher: ACM Press
Full text available:  [pdf\(518.63 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Most commercial database systems do (or should) exploit many sorting techniques that are publicly known, but not readily available in the research literature. These techniques improve both sort performance on modern computer systems and the ability to adapt gracefully to resource fluctuations in multiuser operations. This survey collects many of these techniques for easy reference by students, researchers, and product developers. It covers in-memory sorting, disk-based external sorting, and cons ...

Keywords: Key normalization, asynchronous read-ahead, compression, dynamic memory resource allocation, forecasting, graceful degradation, index operations, key conditioning, nested iteration

Results 1 - 20 of 33

Result page: **1** [2](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

[Go to Google Home](#)

[Web](#) [Images](#) [Video](#) [News](#) [Maps](#) [more »](#)

estimate OR estimating cardinality "data ske

[Search](#)

[Advanced Search](#)
[Preferences](#)

Web Results 1 - 100 of about 274 for **estimate OR estimating cardinality "data skew" formula**. (0.55 seconds)

Solutions Journal: March 99

As we will explain later, **data skew** tends to be an issue more with ad hoc query ... The **formula 1/cardinality** is fine for values which occur with close to ...

www.idug.org/member/journal/mar99/improving_db2.html - 29k -

[Cached](#) - [Similar pages](#)

Solutions Journal: Tuning SQL with Statistics - August 99

Frequency statistics are helpful at identifying **data skew**, which occurs when some column values appear more ... **Estimate** concatenated column **cardinality** ...

www.idug.org/member/journal/aug99/feature1.html - 18k - [Cached](#) - [Similar pages](#)

Inside the Oracle Optimizer, Part 2

To minimize intermediate results, the optimizer attempts to **estimate** the **cardinality** of each result set during the parse phase of SQL execution. ...

www.oracle.com/technology/oramag/webcolumns/

2003/techarticles/burleson_cbo_pt2_pt1.html - 44k - [Cached](#) - [Similar pages](#)

[PDF] Automated design of multi-dimensional clustering tables in ...

File Format: PDF/Adobe Acrobat - [View as HTML](#)

2.3 **Estimating** the **cardinality** of distinct values in a set from a data sample. The ability to **estimate** the **cardinality** of a set from a ...

www.vldb.org/conf/2004/IND5P5.PDF - [Similar pages](#)

[PDF] Estimation of Query-Result Distribution and its Application in ...

File Format: PDF/Adobe Acrobat - [View as HTML](#)

sive for data with high **cardinality**. Hence, most commercial database systems maintain some ... to **estimate** the cost **formula** most accurately, we need rea- ...

www.vldb.org/conf/1996/P448.PDF - [Similar pages](#)

[[More results from www.vldb.org](#)]

[PDF] Data reduction by partial preaggregation - Data Engineering, 2002 ...

File Format: PDF/Adobe Acrobat

Estimating output size. To incorporate partial preaggregation into query optimization we must be able to **estimate** the **cardinality** of the output. ...

ieeexplore.ieee.org/iel5/7807/21451/00994787.pdf - [Similar pages](#)

[PDF] Building large ROLAP data cubes in parallel - Database Engineering ...

File Format: PDF/Adobe Acrobat

sizes and **data skew**. For a fact table with 16 million rows ... we **estimate** val- ... The **cardinality** of the dimensions in a data set can affect ...

ieeexplore.ieee.org/iel5/9219/29237/01319810.pdf - [Similar pages](#)

[[More results from ieeexplore.ieee.org](#)]

[PDF] A Parallel Hash Join Algorithm for Managing Data Skew

File Format: PDF/Adobe Acrobat

Also, the **estimate** of the **cardinality** of the large skew value (shown in black) in the ... timately in **estimating** the subtask time for the composite hash ...

doi.ieeecomputersociety.org/10.1109/71.250117 - [Similar pages](#)

[PDF] A Parallel Sort Merge Join Algorithm for Managing Data Skew

File Format: PDF/Adobe Acrobat

the **cardinality** of the sorted run of relation T at processor p_i and Q_{i,p,r} ... a uniform distribution within a range in **estimating** the time for ...

doi.ieeecomputersociety.org/10.1109/71.205654 - [Similar pages](#)
[[More results from doi.ieeecomputersociety.org](#)]

Apparatus and method for estimating cardinality when data skew is ...

The improved **cardinality estimate** may then be used to. ... Apparatus and method for **estimating cardinality** when **data skew** is present patent application. ...

www.freshpatents.com/

Apparatus-and-method-for-estimating-cardinality-when-data-skew-is-present-dt20050616...

- 20k - [Cached](#) - [Similar pages](#)

Method and apparatus for predicting selectivity of database query ...

The **estimate** of the size of the result set might be performed by the optimizer or some other facility in the database management system. **Estimating** the ...

www.freshpatents.com/

Method-and-apparatus-for-predicting-selectivity-of-database-query-join-conditions-us...

- 92k - Supplemental Result - [Cached](#) - [Similar pages](#)

[[More results from www.freshpatents.com](#)]

Apparatus and method for estimating cardinality when data skew is ...

Without a way to **estimate cardinality** in a manner that accounts for **data skew**, the computer industry will continue to suffer from inaccurate estimates of ...

www.patentdebate.com/PATAPP/20050131914 - Supplemental Result - [Similar pages](#)

[PDF] LEO – DB2's LEarning Optimizer

File Format: PDF/Adobe Acrobat - [View as HTML](#)

errors is the **cardinality estimate** on which costs depend. Cost estimates might be off by 10 ... columns in stored tables, for **estimating** the selectivity of ...

www.dia.uniroma3.it/~vldbproc/006_019.pdf - [Similar pages](#)

[PDF] Distinct Sampling for Highly-Accurate Answers to Distinct Values ...

File Format: PDF/Adobe Acrobat - [View as HTML](#)

in order to **estimate** the number of distinct destination IP. addresses, source-destination pairs, requested urls, etc. **Estimating** the number of distinct ...

www.dia.uniroma3.it/~vldbproc/058_541.pdf - [Similar pages](#)

US Pregrant 20050131914 - Apparatus and method for estimating ...

Apparatus and method for **estimating cardinality** when **data skew** is present ... **cardinality** for a database column using a **formula** that accounts for skew in ...

cxp.paterra.com/uspregrant20050131914.html - 11k - Supplemental Result -

[Cached](#) - [Similar pages](#)

[PDF] Towards Estimating the Number of Distinct Value Combinations for a ...

File Format: PDF/Adobe Acrobat - [View as HTML](#)

cardinality estimation of various relational operators. For example,. the textbook **formula** (which is also widely used in commercial sys- ...

www.cs.toronto.edu/~xhyu/research/papers/cikm05.pdf - [Similar pages](#)

Cost Model. A Mixed Transaction Cost Model. Unwalla

5.4 Partition **cardinality** does not affect migration rate ... Similarly, if there is **data skew** and even if all partitions contain the same number of records, ...

www.techscribe.co.uk/thesis/chapter5.html - 50k - [Cached](#) - [Similar pages](#)

Partitioning and placement. A Mixed Transaction Cost Model. Unwalla

... the possibility that **data skew** makes it impossible to balance partition **cardinality**. ...

Using the **formula** derived therein for **estimating** the number of ...

www.techscribe.co.uk/thesis/chapter2.html - 57k - [Cached](#) - [Similar pages](#)

[PDF] Limiting Result Cardinalities for Multidatabase Queries Using ...

File Format: PDF/Adobe Acrobat

be adjusted by **estimating** the **cardinality** of the sub-query, ... we have only to **estimate** sel ... overcome **data skew** and problems with join sampling. ...

www.springerlink.com/index/DAK9WCYTB42G3BWQ.pdf - [Similar pages](#)

[PDF] On applying hash filters to improving the execution of multi-join ...

File Format: PDF/Adobe Acrobat

of **estimating the cardinality** of R. j. projected on the non- ... mance model of data **skew** effects in parallel joins. Proc 17th Int Conf ...

www.springerlink.com/index/8YFXC93B2CM28N3M.pdf - [Similar pages](#)

[[More results from www.springerlink.com](#)]

[PS] Probabilistic Optimization of Top N Queries

File Format: Adobe PostScript - [View as Text](#)

4 **Estimating** Initial Probability Densi-. ties. We have discussed how to propagate **cardinality** den-. sities through the plan tree, by multiplying the op- ...

www.cs.wisc.edu/~donjerko/topN.ps - [Similar pages](#)

Citations: Statistical profile estimation in database systems ...

Several techniques have been proposed in the literature to **estimate** ... On **Estimating the Cardinality** of Aggregate Views - Paolo Ciaccia Matteo (Correct) ...

citeseer.ist.psu.edu/context/2589/0 - 63k - [Cached](#) - [Similar pages](#)

Citations: Sequential sampling procedures for query size ...

Semantic **Cardinality** Estimation for Mapping and.. - Nodine, Cherniack. ... and data **skew**, we only have to modify the **formula** for **estimating the** ...

citeseer.ist.psu.edu/context/15779/0 - 35k - [Cached](#) - [Similar pages](#)

[PS] ON APPLYING HASH FILTERS TO IMPROVING THE EXECUTION OF MULTI-JOIN ...

File Format: Adobe PostScript - [View as Text](#)

only have to modify the corresponding **formula** accordingly [10]. ... has two attributes A and B. The problem of **estimating the cardinality** of R ...

www.ee.ntu.edu.tw/~mschen/paperps/hfp.ps - [Similar pages](#)

[PS] APPLYING HASH FILTERS TO IMPROVING THE EXECUTION OF BUSHY TREES

File Format: Adobe PostScript - [View as Text](#)

In the presence of **data skew** [26], we only have to. modify the corresponding **formula** accordingly [9]. ... problem of **estimating the cardinality** of R ...

www.ee.ntu.edu.tw/~mschen/paperps/hfvldb93.ps - [Similar pages](#)

[[More results from www.ee.ntu.edu.tw](#)]

[PDF] Limiting Result Cardinalities for Multidatabase Queries using ...

File Format: PDF/Adobe Acrobat - [View as HTML](#)

This requires **estimating the stop. cardinality** by using database statistics as well as a ... las presented in [PSC84] we can **estimate the selectivity** ...

fusion.cs.uni-magdeburg.de/pubs/bncod01.pdf - [Similar pages](#)

[PDF] On Incorporating Iceberg Queries in Query Processors

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Dataset refers to the name of the dataset, **Cardinality** indicates the number of ... between SMA and CIQE based on the conditional in the **formula** on line ...

dsl.serc.iisc.ernet.in/publications/conference/iceberg.pdf - [Similar pages](#)

[PDF] j ¢£j ¢ ¥ ¥ ¢ ¢£j © ¢ ¢ ¥ © ¥ ¢ ¢£j ¥ ¥ ¢ ¥

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Or, at what is the minimum **data skew** factor for CIQE to be effective for a ...

Estimating the. ¢"A#¢\$}. mentioned in this. **formula** is simple and is done the ...

dsl.serc.iisc.ernet.in/publications/report/TR/TR-2002-01.pdf - [Similar pages](#)

[[More results from dsl.serc.iisc.ernet.in](#)]

Gathering Optimizer Statistics

If the optimizer determines that the selectivity, cost, and **cardinality** ... Histograms provide improved selectivity estimates in the presence of **data skew**, ...

www.utexas.edu/its/unix/reference/oracledocs/v92/B10501_01/server.920/a96533/stats.htm - 93k -

[Cached](#) - [Similar pages](#)

[PDF] [Efficient Computation of Multiple Group By Queries](#)

File Format: PDF/Adobe Acrobat

able to **estimate** the **cardinality** of a Group By query, which is a. hard problem. For this, we assume that known techniques for. **estimating** number of distinct ...
portal.acm.org/ft_gateway.cfm?id=1066188&type=pdf - [Similar pages](#)

[PDF] [Dynamic query scheduling in parallel data warehouses](#)

File Format: PDF/Adobe Acrobat

standard bitmaps for the low-**cardinality** dimensions Time and Channel, ... Handling **data skew** in parallel hash join computation using two-phase scheduling. ...
<dx.doi.org/10.1002/cpe.786> - [Similar pages](#)

[PDF] [Applying Hash Filters to Improving the Execution of Bushy Trees](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

In the presence of **data skew** [26], we only have to. modify the corresponding formula accordingly [S]. ... problem of **estimating** the **cardinality** of ...
<www.sigmod.org/vldb/conf/1993/P505.PDF> - [Similar pages](#)

[PDF] [A Taxonomy and Performance Model of Data Skew Effects in Parallel ...](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

cardinality and memory size. These scalar parameters. are adequate to model a single node system or a parallel. system with no **data skew**. ...
<www.sigmod.org/vldb/conf/1991/P537.PDF> - [Similar pages](#)
[[More results from www.sigmod.org](#)]

[PDF] [Building Large ROLAP Data Cubes in Parallel](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

and assigns to every view identifier an **estimate** of the size of the respective view [9, ...
The **cardinality** of the dimensions in a data set can affect the ...
projects.cs.dal.ca/panda/papers/ideas_04.pdf - [Similar pages](#)

[PDF] [Selectivity Estimation of Complex Spatial Queries](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

w (output **cardinality** - OC) is estimated by the following formula: ... Belussi A., Faloutsos C. **Estimating** the Selectivity of Spatial Queries Using the ...
<www.cwi.nl/htbin/ins1/publications?request=pdf&key=MaPa:SSTD:01> - [Similar pages](#)

[PDF] [stefan.manegold, florian.waas @cwi.nl Parallel database research ...](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

forms in the presence of highly skewed **data**. **Skew** in general means that due to the ...
The first series of experiments gives an overall **estimate** for the ...
<www.cwi.nl/htbin/ins1/publications?request=pdf&key=MaWa:HPCS:98> - [Similar pages](#)

[PDF] [Improved Data Partitioning For Building Large ROLAP Data Cubes in ...](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

speedup, scaleup, sizeup, output sizes and **data skew**. ... **cardinality** of some dimensions may be small, such as gender, ... **estimate** values n ...
<users.cs.dal.ca/~arc/publications/1-20/paper.pdf> - [Similar pages](#)

[Märtens, Holger; Rahm, Erhard; Stöhr, Thomas: Dynamic Query ...](#)

We employ standard bitmaps for the low-**cardinality** dimensions Time and Channel, ...
A Taxonomy and Performance Model of **Data Skew** Effects in Parallel Joins. ...
<lips.informatik.uni-leipzig.de/pub/2001-38/en> - 61k - [Cached](#) - [Similar pages](#)

[Stöhr, Thomas; Märtens, Holger; Rahm, Erhard: Multi-Dimensional ...](#)

Simple bitmap indices become inefficient for high-**cardinality** attributes resulting in ...
data skew effects as well as the consequences of multi-user mode. ...
<lips.informatik.uni-leipzig.de/pub/2000-12> - 70k - [Cached](#) - [Similar pages](#)
[[More results from lips.informatik.uni-leipzig.de](#)]

[PDF] [1 Introduction](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

A new **estimate** for the **cardinality** of the result. can be made by multiplying the monitored ... For **estimating** the data rates there are two possible ways. ...
www.ncl.ac.uk/polarstar/monitoringTR.pdf - [Similar pages](#)

[PDF] [Automating Physical Database Design in a Parallel Database](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

statistics on the database, including table **cardinality**, ... scope of this paper. Avoid **Data Skew**. When **estimating** new statis- ...

www.almaden.ibm.com/software/projects/autonomic/partitioning_advisor.pdf - [Similar pages](#)

[PDF] [Categories and Subject Descriptors Keywords](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

is resistant to **data skew**. It can also join multiple relations while still ... particular when the **cardinality** of the join index is high. ...

www.cs.brown.edu/courses/cs227/Papers/Indexing/Ross97.pdf - [Similar pages](#)

[Learning from empirical results in query optimization - Patent 6763359](#)

R. Ahad, K.V.B. Rao, and D. McLeod, On **Estimating the Cardinality** of the ... success in using histograms to detect and correct for **data skew** [IC91, PIHS96, ...

www.freepatentsonline.com/6763359.html - 90k - [Cached](#) - [Similar pages](#)

[Joining two database relations on a common field in a parallel ...](#)

(o) said third means for **estimating**, means for subpartitioning and third means for updating iteratively operating until there is no execution time **estimate** ...

www.freepatentsonline.com/5121494.html - 79k - [Cached](#) - [Similar pages](#)

[[More results from www.freepatentsonline.com](#)]

[PDF] [Dynamic Load Balancing in Hierarchical Parallel Database Systems](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

a **cardinality** is randomly chosen in one of the small, medium or large ranges. ... is severe **data skew** (which yields high variations in ...

www-smis.inria.fr/~bouganim/Publis/BOUGA_CI3_VLDB_1996.pdf - [Similar pages](#)

[PDF] [IDEA Parallel Execution Model for IDEA](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

in the presence of **data skew** (see Section 5.4); each thread chooses the ... lots of tuples. a is then equal to the **cardinality** of the pipelined collection. ...

www-smis.inria.fr/~bouganim/Publis/BOUGA_RE5_IDEA_1995.pdf - [Similar pages](#)

[Using Parallel Execution](#)

Consider **data skew**. If a join key involves excessive **data skew**, ... The sample **cardinality** can also be used, in some cases, to **estimate** table **cardinality**. ...

www.lc.leidenuniv.nl/awcourse/oracle/server.920/a96520/tuningpe.htm - 315k -

[Cached](#) - [Similar pages](#)

[PDF] [LEO – DB2's LEarning Optimizer](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

modeling errors is the **cardinality estimate** on which costs depend. ... success in using histograms to detect and correct for **data skew** [IC91, PIHS96, PI97], ...

<https://www.almaden.ibm.com/software//projects/autonomic/LeoVLDB01.pdf> -

[Similar pages](#)

[PDF] ["Query Cardinality Estimation in Relational Databases." Irena ...](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

data flow of each node, it is necessary to **estimate** the sizes of its ... stage, while **estimating** its **cardinality**, in order to make sampling less expensive. ...

www.cse.yorku.ca/~irka/last/thesis.pdf - Supplemental Result - [Similar pages](#)

[PDF] [A Survey of Query Optimization in Parallel Databases Technical ...](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Given the **cardinality** of the Age attribute and declustering information, ... we monitor

the effect of **data skew**, the effect of I/O bandwidth and the ...
www.cse.yorku.ca/techreports/1999/CS-1999-04.pdf - [Similar pages](#)
[[More results from www.cse.yorku.ca](#)]

[PPT] [Sabre Lab Reporting](#)

File Format: Microsoft Powerpoint - [View as HTML](#)

But sensitive to **data skew**. Sort-merge is better when one or both inputs ... Number of tuples (**cardinality**) and number of blocks for each table and index. ...
lambda.uta.edu/cse6331/spring03/l1.ppt - [Similar pages](#)

[PDF] [Sabre Lab Reporting](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

To **estimate** the cost of a plan operator (such as block-nested ... Number of tuples (**cardinality**) and number of blocks for each table and ...
lambda.uta.edu/cse6331/fall04/l1.pdf - [Similar pages](#)

[Gathering Statistics](#)

The optimizer uses these statistics to **estimate** how much I/O and memory ...
Histograms provide improved selectivity estimates in the presence of **data skew**, ...
www.mid.main.vsu.ru/docs/oracle/server.816/a76992/stats.htm - 77k -
[Cached](#) - [Similar pages](#)

[Method of generating attribute **cardinality** maps - usptntsrv](#)

From the above cost **formula**, it is apparent that the following quantities are crucial for the query optimizer in order to **estimate** the cost of such an ...
www.uspatentserver.com/686/6865567.html - 120k - Supplemental Result -
[Cached](#) - [Similar pages](#)

[Gathering Statistics](#)

The optimizer uses these statistics to **estimate** how much I/O and memory are ... to get a representative **Cardinality** estimation for the given expression. ...
www.mcs.csu Hayward.edu/support/oracle/doc/8.1.7/server.817/a76992/stats.htm - 77k -
[Cached](#) - [Similar pages](#)

[PDF] [Complex Spatial Query Processing*](#)

File Format: PDF/Adobe Acrobat

The current **formula** is accurate only for the initial version of the algorithm (ST_SS).
We can easily adapt it to **estimate** the number of problems run by ...
www.ingentaconnect.com/content/klu/gein/2004/00000008/00000004/05277173 -
[Similar pages](#)

[PDF] [SPADE: An Efficient Algorithm for Mining Frequent Sequences](#)

File Format: PDF/Adobe Acrobat

of sampling, can be sensitive to the **data-skew**, which can adversely affect ...
cardinality of L(Z), denotes the number of distinct sid values in the id-list ...
www.ingentaconnect.com/content/klu/mach/2001/00000042/F0020001/00279319 -
[Similar pages](#)

[PDF] [Generalized Multi-dimensional Data Mapping and Query Processing](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

which can be optimized by **estimating** the final search radius. ... in fact shifts the position of the line $y = 1-x$ so that it can adapt to the **data skew**. ...
www.comp.nus.edu.sg/~kalnis/tods05.pdf - [Similar pages](#)

[PDF] [Performance Enhancements for Hybrid Hash Join](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

The dominant parts in this **formula** are the number of recursion levels L and input ... In complex queries, it is impossible to **estimate** sizes of intermediate ...
www.comp.nus.edu.sg/~caoyu/papers/Performance%20Enhancements%20for%20Hybrid%20Hash%20Join.pdf -
Supplemental Result - [Similar pages](#)

[PDF] [Algorithms and frameworks for stream mining and knowledge ...](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

due to **data skew** present in the various dataset partitions [35]. ... Table 3 illustrates these datasets along with their **cardinality**. ...

www.coregrid.net/mambo/images/stories/TechnicalReports/tr-0046.pdf - [Similar pages](#)

[PDF] [Processing Queries for First-Few Answers](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

form of **cardinality** constraints (such as existence and func- ... sary for **estimating** the result of a semi-join is the semi-join. selectivity of a predicate, ...

www.bayardo.org/ps/cikm96.pdf - [Similar pages](#)

[PDF] [IBM Research Report](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Inverse Cardenas' formula Car75 to rst get the per- ... scope of this paper. Avoid **Data Skew**:. When **estimating** new statis- ...

[domino.watson.ibm.com/library/CyberDig.nsf/papers/D6D46969B5619FE985256BA00058849F/\\$File/rj10225.pdf](http://domino.watson.ibm.com/library/CyberDig.nsf/papers/D6D46969B5619FE985256BA00058849F/$File/rj10225.pdf) - Supplemental

Result - [Similar pages](#)

[PDF] [Distributed Approximate Mining of Frequent Patterns](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Distributed One-pass Partition when the **data skew** between. the data partitions is high. ... of these datasets along with their **cardinality**. Each dataset ...

www.dsi.unive.it/~claudio/papers/AP_ACMSAC05.pdf - [Similar pages](#)

[PDF] [Approximate Mining of Frequent Patterns on Streams](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

a lot of false positives due to **data skew** present in the various dataset partitions ...

Table 3 illustrates these datasets along with their **cardinality**. ...

www.dsi.unive.it/~claudio/papers/AP_stream_IDA_journal.pdf - [Similar pages](#)

[Gathering Optimizer Statistics](#)

The optimizer uses the selectivity of a predicate to **estimate** the cost ... Histograms provide improved selectivity estimates in the presence of **data skew**, ...

www.cs.uvm.edu/oracle9doc/server.901/a87503/stats.htm - 91k -

[Cached](#) - [Similar pages](#)

[PDF] [Wavelet Synopses with Error Guarantees](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

value of **data skew**, our MinRelVar and MinRelBias algorithms are ... tributes have low **cardinality**, but there are 10 quantitative attributes, ...

infolab.usc.edu/csci599/Fall2002/paper/VL.pdf - [Similar pages](#)

[Gathering Optimizer Statistics](#)

The optimizer uses the selectivity of a predicate to **estimate** the cost of a ... cost , and **cardinality** of two finalist indexes is the same , then it uses ...

barry1.cityu.edu.hk/han3/2/6/9/1/1/1/1/www.cs.cityu.edu.hk/cslab/manuals/oracle9i/server.901/a87503/stats.htm - 104k - Supplemental Result -

[Cached](#) - [Similar pages](#)

[PDF] [An Improved Algorithm for Implication Testing Involving Arithmetic ...](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

The reason is that the **estimate** of the task time for HJ* becomes very. difficult: A high skew value for one relation the join output **cardinality** ...

lib.nau.edu.ua/acm/disk2/text/3-3/TKDE6/P0997.pdf - Supplemental Result -

[Similar pages](#)

[PDF] [Competitive Online Comparison for Parallel Joins](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

A critical point here is to **estimate** the task sizes that will be distributed to the processing ... tribute values, using a Zipf-based formula [23]: ...

lib.nau.edu.ua/acm/2003/out/papers/sbbd/sergio.pdf - Supplemental Result -

[Similar pages](#)

[[More results from lib.nau.edu.ua](#)]

[PDF] [Microsoft PowerPoint - DBCC 302 310 and friends.ppt](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Ensure proper Selectivity **formula** is being used ... **Estimate**: indid 1, selectivity

0.100000, rows 1 pages 3. index height 1. Run UPDATE STATISTICS ...

www.benstopford.com/software/dbcc%20traceon%20guide.pdf - Supplemental Result -

[Similar pages](#)

[Designing multidimensional clustering \(MDC\) tables](#)

7 Also, note that on average (assuming little or no **data skew**), ... one or more dimensions to a coarser granularity in order to give it a lower **cardinality**. ...

publib.boulder.ibm.com/infocenter/

db2help/topic/com.ibm.db2.udb.doc/admin/c0007238.htm - 26k -

[Cached](#) - [Similar pages](#)

[PDF] [Distributed and Stream Data Mining Algorithms for Frequent Pattern ...](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

with their **cardinality** and the minimum and maximum sizes of the ... mining on streams when the **data skew** between different incoming blocks is high. ...

www.unive.it/media/diplInformatica/phd/TD-2006-4.pdf - [Similar pages](#)

[PDF] [-Diversity: Privacy Beyond k-Anonymity](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

The large groups were caused by **data skew**. For example, there were only 114 tuples

... Principle Component Analysis (PCA) and the Bayes **Estimate** (BE) to ...

[www.cs.cornell.edu/people/dkifer/diversityTKDDraft.pdf](http://www.cs.cornell.edu/people/dkifer/diversityTKDDdraft.pdf) - [Similar pages](#)

[PS] [-Diversity: Privacy Beyond k-Anonymity Ashwin Machanavajjhala ...](#)

File Format: Adobe PostScript - [View as HTML](#)

Given f and T ?, we can derive using bayesian reasoninga **formula** for. $f(q,s,T?)$. This is a key technical result ofthis paper since in this **formula** lies ...

www.cs.cornell.edu/people/dkifer/papers/diversitylong.ps - Supplemental Result -

[Similar pages](#)

[Tuning Parallel Execution](#)

This value is INSTANCES in the **formula**. **Estimate** the maximum number of ...

Furthermore, this can lead to **data skew** during the execution if you do not ...

www.csee.umbc.edu/help/oracle8/server.815/a67775/ch24_pex.htm - 253k -

[Cached](#) - [Similar pages](#)

[24 Using Parallel Execution](#)

Now, assume that you have a join key with high **cardinality**, but one of the values ...

Consider **data skew**. If a join key involves excessive **data skew**, ...

www.stanford.edu/dept/itss/docs/oracle/10g/server.101/b10736/usingpe.htm - 273k -

[Cached](#) - [Similar pages](#)

[PDF] [Oracle9i Data Warehousing Guide](#)

File Format: PDF/Adobe Acrobat

To avoid **data skew** among partitions. Hash partitioning is an effective means of ...

ESTIMATE STATISTICS to collect **cardinality** statistics on all fact ...

www.indiana.edu/~dss/Services/Documentation/DW_GuideNotes.pdf - [Similar pages](#)

[PS] [Wavelet Synopses with Error Guarantees](#)

File Format: Adobe PostScript - [View as Text](#)

the **estimate** has over 2000% relative error! Likewise, d ... value of **data skew**, our MinRelVar and MinRelBias algorithms are ...

www.pittsburgh.intel-research.net/people/gibbons/papers/wavelets-sigmod02.ps -

[Similar pages](#)

[PDF] [A Survey of Query Optimization in Parallel Databases Technical ...](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Given the **cardinality** of the Age attribute and declustering information, ... This **formula** is considered as a general approximation for joining large ...

miami.int.gu.edu.au/dbs/papers/Other/parallel-query.pdf - Supplemental Result -

[Similar pages](#)

[PDF] [ABSTRACT](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

data skew problems. We selected maxdiff histogram method with sampling by ... selectivities for sample values and then **estimate cardinality** for whole table. ...

www.lib.ncsu.edu/theses/available/etd-01052005-235202/unrestricted/etd.pdf -

[Similar pages](#)

[PDF] [arXiv:cs.DB/0501029 v1 14 Jan 2005](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

taken from the set $\{1, \dots, y\}$ and having **cardinality** z : ... **formula** for s it follows that $m \cdot b = s$, thus $m = s/b$. Consider now the vector ...

132.236.180.11/pdf/cs.DB/0501029 - Supplemental Result - [Similar pages](#)

[Method of generating attribute **cardinality** maps Number:6865567 ...](#)

b. **estimating** a value associated with an element based on the location of the straight line at the element; and c. using the **estimate** for determining an ...

www.linkgrinder.com/Patents/Method_of_generating_attribute_cardin_6865567.html -

118k - Supplemental Result - [Cached](#) - [Similar pages](#)

[PDF] [Building Large ROLAP Data Cubes in Parallel](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

sizes and **data skew**. For a fact table with 16 million rows ... signs to every view identifier an **estimate** of the size of the. respective view [9, 23]. ...

134.117.20.125/www/publications/2-73.pdf - Supplemental Result - [Similar pages](#)

[PDF] [Improved Data Partitioning For Building Large ROLAP Data Cubes in ...](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

every view identifier an **estimate** of the size of the respective view (Flajolet & Martin, ... **cardinality** of some dimensions may be small, such as gender, ...

134.117.20.125/www/publications/1-41.pdf - Supplemental Result - [Similar pages](#)

[PDF] [Meet the HP Superdome servers](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Cardinality is the number of unique values that can appear for a particular ... Make sure the catalog statistics account for **data skew**, because skewed data ...

h71028.www7.hp.com/ERC/downloads/NSSQLbest%20practicesWP.pdf -

[Similar pages](#)

[PDF] [Integration of Spatial Join Algorithms for Processing Multiple Inputs](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

formula that predicts the cost of RJ in terms of node accesses,. based on the properties (density, **cardinality**) of the joined data- ...

acm.org/sigmod/sigmod99/e proceedings/papers/papadias.pdf - [Similar pages](#)

[PDF] [Efficient Processing of Joins on Set-valued Attributes](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

We can use this **formula** to **estimate** the size of S ... structures for set-valued attributes of low **cardinality**. In Technical Report, University of Mannheim, ...

acm.org/sigmod/sigmod03/e proceedings/papers/r06p02.pdf - [Similar pages](#)

[PDF] [ABSTRACT ACCURATE DATA APPROXIMATION IN CONSTRAINED ENVIRONMENTS ...](#)

File Format: PDF/Adobe Acrobat

algorithm relies on three basic operations: (1) **Estimating** the ... produces tight error bounds for such a large **data skew** value and for small data ...

https://drum.umd.edu/dspace/bitstream/1903/2681/1/umi-umd-2601.pdf - [Similar pages](#)

[PDF] Indexing and Statistics Strategies for DB2 UDB for iSeries

File Format: PDF/Adobe Acrobat - [View as HTML](#)

application development world, there is not a simple formula for the ... information about data skew and column cardinality becomes critically important ...

www-03.ibm.com/servers/enable/site/bi/strategy/strategy.pdf - [Similar pages](#)

[PDF] HKUST Institutional Repository

File Format: PDF/Adobe Acrobat

w (output cardinality - OC) is estimated by the following formula: ... lems due to data skew. Another direction for future work is the study of alternative ...

repository.ust.hk/dspace/bitstream/1783.1/2517/1/sele.pdf - [Similar pages](#)

[PS] SPADE: An Ecient Algorithm for Mining Frequent Sequences

File Format: Adobe PostScript - [View as Text](#)

from the formula above, but a closed-form expression for an upper-bound on the ... A simple check on the cardinality (i.e., the unique sid values) ...

www.cs.helsinki.fi/u/gionis/seminar_papers/zaki00spade.ps - [Similar pages](#)

[PDF] SPADE: An Efficient Algorithm for Mining Frequent Sequences

File Format: PDF/Adobe Acrobat - [View as HTML](#)

from the formula above, but a closed-form expression for an upper-bound on the ... cardinality of $L(Z)$, denotes the number of distinct sid values in the ...

www.cs.rpi.edu/~zaki/PS/MLJ01.pdf - Supplemental Result - [Similar pages](#)

[PS] Bulletin of the Technical Committee on March 1998 Vol. 21 No. 1 ...

File Format: Adobe PostScript - [View as Text](#)

C and X. However, this joint density is rarely known and very difficult to estimate.

Hence, one has to resort to. various techniques for estimating. ...

sites.computer.org/debull/mar98-a4draft.ps - [Similar pages](#)

[PDF] Bulletin of the Technical Committee on March 1998 Vol. 21 No. 1 ...

File Format: PDF/Adobe Acrobat - [View as HTML](#)

However, this joint density is rarely known and very difficult to estimate. Hence, one has to resort to. various techniques for estimating. ...

sites.computer.org/debull/98MAR-CD.pdf - [Similar pages](#)

[PDF] Powerful Database Support for High Performance Data Mining

File Format: PDF/Adobe Acrobat

object replication. • and data skew. for the estimation which clearly limits the value of this formula. 3.2.5 Join Algorithms Based on Sorting ...

www.dbs.informatik.uni-muenchen.de/~boehm/publications/Habil.pdf - [Similar pages](#)

[PDF] 07 chap

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Estimating the cost and the output size of a spatial join is an important and difficult. problem. ... Data skew is the most difficult factor to estimate, ...

isl.cs.unipi.gr/db/pubs/chapters/sdb05.pdf - [Similar pages](#)

AskTom "Hints When to use"

Also estimating stats are not 100% secure for some 100% important SQL's. ... I use the cardinality hint mostly with the where x in (select from ...

asktom.oracle.com/.../f?p=4950:8:2177642270773127589::NO::F4950_

P8_DISPLAYID,F4950_P8_CRITERIA:7038986332061 - 218k -

[Cached](#) - [Similar pages](#)

AskTom "Hints When to use"

Also estimating stats are not 100% secure for some 100% important SQL's. ... I was surprised to see cardinality listed in the good hint list. ...

asktom.oracle.com/pls/ask/

f?p=4950:8:2177642270773127589::NO::LAST_PAGE:8:YES - 209k -

[Cached](#) - [Similar pages](#)

Tuning Parallel Execution

For most installations , this number will be 1 . This value is INSTANCES in the formula

. **Estimate** the maximum number of concurrent queries executing at ...

cs.mipt.ru/docs/comp/eng/develop/databases/

oracle/8i/doc/server.815/a67775/ch24_pex.htm - 247k - Supplemental Result -

Cached - Similar pages

[PDF] Query Evaluation Techniques for Large Databases

File Format: PDF/Adobe Acrobat - View as HTML

Both issues require adjusting the fan-in of the first merge step using the formula given









below,. Query Evaluation Survey. 12. Draft of Friday, June 5, 1992, ...

www.cse.iitb.ac.in/maloo/ddp/queryeval-large-db.pdf - Supplemental Result -

Similar pages

Result Page: 1 2 **Next**


Free! Get the Google Toolbar. Download Now - About Toolbar

Google	<input type="text"/>			Search		PageRank		Check		AutoLink		AutoFill		Options		Highlight
--------	----------------------	---	---	--------	---	----------	---	-------	---	----------	---	----------	---	---------	---	-----------

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Google Home - Advertising Programs - Business Solutions - About Google

©2006 Google

 | [Search History](#) | [My Account](#) | [Sign out](#)

[Go to Google Home](#)

[Web](#) [Images](#) [Video](#) [News](#) [Maps](#) [more »](#)

[Advanced Search](#)
[Preferences](#)

Web

Results 1 - 7 of about 11 for "**estimating cardinality**" "**data skew**". (0.28 seconds)

Apparatus and method for estimating cardinality when data skew is ...

A cardinality estimator computes an estimated cardinality for a database column using a formula that accounts for skew in the data in the column.

www.patentdebate.com/PATAPP/20050131914 - Supplemental Result - [Similar pages](#)

US Pregrant 20050131914 - Apparatus and method for estimating ...

US Pregrant Abstract 20050131914 (Full Text at USPTO) , 2005 -06 -16 (Class /Claims View). Apparatus and method for **estimating cardinality when data skew** ...
cxp.paterra.com/uspregrant20050131914.html - 11k - Supplemental Result -

[Cached](#) - [Similar pages](#)

美国专利申请公开说明书20050131914 - Apparatus and method for ... - [[Translate this page](#)]

美国专利申请公开说明书20050131914 (美国专利商标局的全文) , 2005-06-16.

Apparatus and method for **estimating cardinality when data skew** is present ...

cxp.paterra.com/uspregrant20050131914cn.html - 18k - Supplemental Result -

[Cached](#) - [Similar pages](#)

[PDF] Estimating Progress of Execution for SQL Queries

File Format: PDF/Adobe Acrobat - [View as HTML](#)

A key challenge arises from **estimating cardinality** of nodes of ... Evaluate robustness of our estimator when **data skew** is. varied. ...

www.cse.iitb.ac.in/dbms/Data/Courses/CS632/Papers/progress-est-surajit04.pdf -

[Similar pages](#)

AskTom "Performance poor after analyzing table"

I ran into a situation with severe **data skew** today this helped me track down the ...

How CBO is **estimating cardinality** for queries using indexes on strings ...

asktom.oracle.com/pls/ask/f?p=4950:8:::::F4950_P8_DISPLAYID:3126073805757 -

316k - [Cached](#) - [Similar pages](#)

Martin & Associates, LLC IBM Intellectual Property Law Department ...

20050131914 - Apparatus and method for **estimating cardinality when data skew** is present 20050125668 - Apparatus and method for inter-program authentication ...

www.freshpatents.com/

Martin-Associates-LLC-IBM-Intellectual-Property-Law-Department-cndirm.php - 14k -

[Cached](#) - [Similar pages](#)

Data processing: database and file management or data structures ...

20050131914 - Apparatus and method for **estimating cardinality when data skew** is present: A cardinality estimator computes an estimated cardinality for a ...

www.freshpatents.com/

Data-processing--database-and-file-management-or-data-structures-dt200506ntc707.php

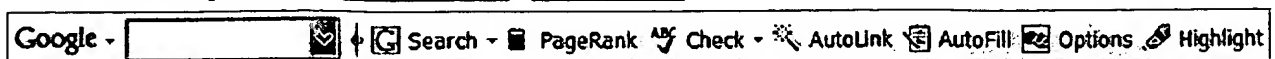
- 257k - [Cached](#) - [Similar pages](#)

[[More results from www.freshpatents.com](#)]

In order to show you the most relevant results, we have omitted some entries very similar to the 7 already displayed.

If you like, you can repeat the search with the omitted results included.

Free! Get the Google Toolbar. [Download Now](#) - [About Toolbar](#)



[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

oogoom@gmail.com | [Search History](#) | [My Account](#) | [Sign out](#)

[Go to Google Home](#)

[Web](#) [Images](#) [Video](#) [News](#) [Maps](#) [more »](#)

"estimating cardinality" "data skew"

[Search](#)

[Advanced Search](#)
[Preferences](#)

Web

Results 1 - 7 of about 11 for "**estimating cardinality**" "**data skew**". (0.28 seconds)

Apparatus and method for estimating cardinality when data skew is ...

A cardinality estimator computes an estimated cardinality for a database column using a formula that accounts for skew in the data in the column.

www.patentdebate.com/PATAPP/20050131914 - Supplemental Result - [Similar pages](#)

US Pregrant 20050131914 - Apparatus and method for estimating ...

US Pregrant Abstract 20050131914 (Full Text at USPTO) , 2005 -06 -16 (Class /Claims View). Apparatus and method for **estimating cardinality** when **data skew** ...
cxp.paterra.com/uspregrant20050131914.html - 11k - Supplemental Result -

[Cached](#) - [Similar pages](#)

美国专利申请公开说明书20050131914 - Apparatus and method for ... - [[Translate this page](#)]

美国专利申请公开说明书20050131914 (美国专利商标局的全文) , 2005-06-16.

Apparatus and method for **estimating cardinality** when **data skew** is present ...

cxp.paterra.com/uspregrant20050131914cn.html - 18k - Supplemental Result -

[Cached](#) - [Similar pages](#)

[PDF] Estimating Progress of Execution for SQL Queries

File Format: PDF/Adobe Acrobat - [View as HTML](#)

A key challenge arises from **estimating cardinality** of nodes of ... Evaluate robustness of our estimator when **data skew** is. varied. ...

www.cse.iitb.ac.in/dbms/Data/Courses/CS632/Papers/progress-est-surajit04.pdf -

[Similar pages](#)

AskTom "Performance poor after analyzing table"

I ran into a situation with severe **data skew** today this helped me track down the ...

How CBO is **estimating cardinality** for queries using indexes on strings ...

asktom.oracle.com/pls/ask/f?p=4950:8:::::F4950_P8_DISPLAYID:3126073805757 -

316k - [Cached](#) - [Similar pages](#)

Martin & Associates, LLC IBM Intellectual Property Law Department ...

20050131914 - Apparatus and method for **estimating cardinality** when **data skew** is present 20050125668 - Apparatus and method for inter-program authentication ...

www.freshpatents.com/

Martin-Associates-LLC-IBM-Intellectual-Property-Law-Department-cndirm.php - 14k -

[Cached](#) - [Similar pages](#)

Data processing: database and file management or data structures ...

20050131914 - Apparatus and method for **estimating cardinality** when **data skew** is present: A cardinality estimator computes an estimated cardinality for a ...

www.freshpatents.com/

Data-processing--database-and-file-management-or-data-structures-dt200506ntc707.php

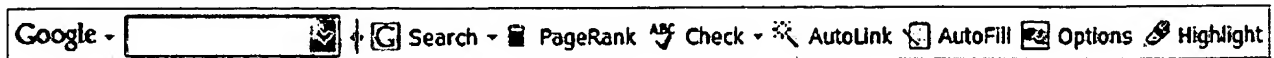
- 257k - [Cached](#) - [Similar pages](#)

[[More results from www.freshpatents.com](#)]

In order to show you the most relevant results, we have omitted some entries very similar to the 7 already displayed.

If you like, you can repeat the search with the omitted results included.

Free! Get the Google Toolbar. [Download Now](#) - [About Toolbar](#)



"estimating cardinality" "data skew"

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google